S&P Global
Market Intelligence
451 Research
Business Impact Brief

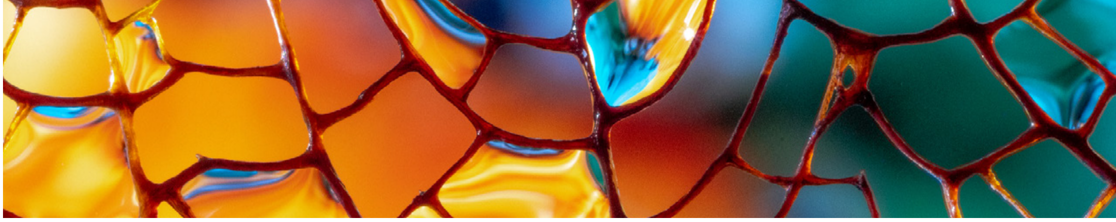

Generative AI: Balancing innovation and privacy

The Take

Artificial intelligence — especially generative AI (GenAI) — has transformative business potential, but also inherent risks. Although the large language models (LLMs) underpinning GenAI are rarely built in-house, other models often are. When AI technology is homegrown, organizations must assume full responsibility for ensuring security, privacy and compliance for their data and activities. Public clouds have their own risks, despite their popularity as an AI deployment venue. In these environments, ownership and control of data or intellectual property (IP) is often in question, and data may be exposed to train models that benefit others.

Figure 1: Security and privacy are top-ranked data challenges in organizational AI efforts

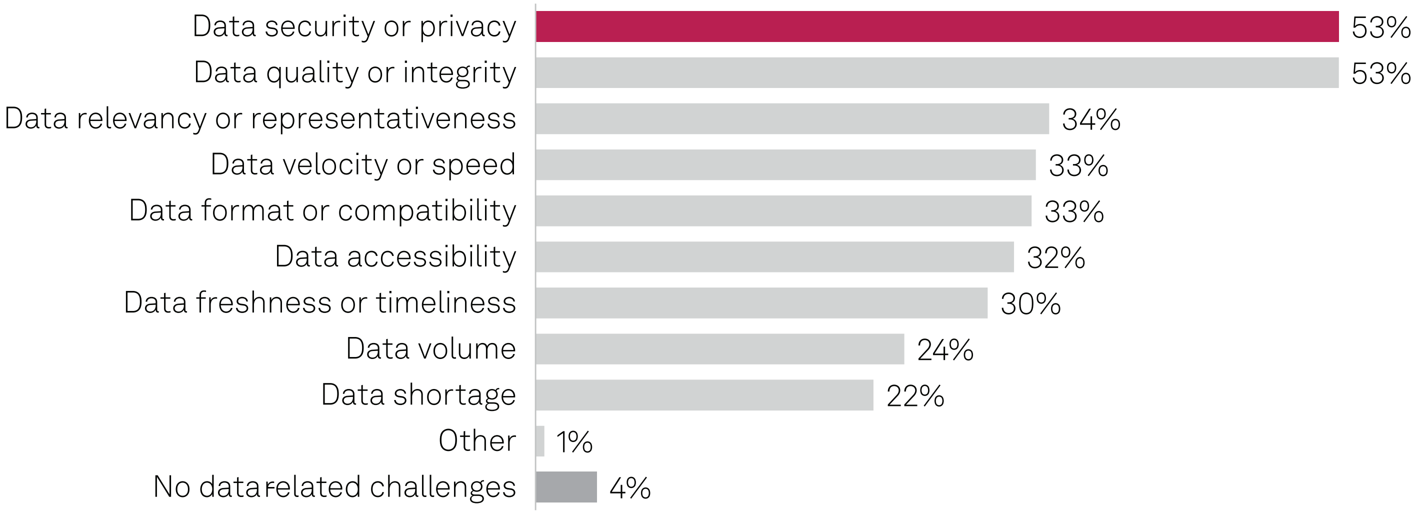

Q. Which of the following data-related challenges are most apparent in your organization's current effort to develop AI models and/or utilize AI-enabled technologies? Please select all that apply.
Base: Organizations that have adopted AI or have AI-related initiatives (n=343).
Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data Architecture for AI 2024.

Deploying GenAI can be particularly difficult for businesses looking to safeguard data assets because it is rare to have full control of underlying LLMs. Sensitive datasets are especially at risk, and organizations are taking steps to increase control, signifying concerns. For example, in our Voice of the Enterprise: Data & Analytics, Data Architecture for AI 2024 survey, 62% of respondents from organizations that primarily buy or acquire their AI technology report using technical controls to protect sensitive organizational data.

Ultimately, privacy and security are top challenges for organizations as they deploy AI and GenAI. Successfully tackling them can mean the difference between AI remaining a risky cost center or becoming a profitable and compliant business transformer. Addressing privacy challenges requires protecting IP, securing personal or sensitive data and controlling access to appropriate AI models, yet much work remains to support the privacy and security of datasets used in AI. Based on the same survey cited above, only 17% of respondents from organizations that have low digital maturity report "very good" privacy and security of datasets used to support AI initiatives. Furthermore, leadership and guidance are often lacking. Only 42% of respondents from organizations with AI initiatives report having a corporate AI ethics board.

AI governance is often proposed as an answer to these challenges by giving structure and accountability to AI processes and technology so efforts are safe and scalable. Yet only a slim majority (52%) of organizations with ongoing AI initiatives report engaging in AI governance practices. There is an enormous opportunity for businesses to bolster AI governance programs. However, AI governance is a multifaceted practice. *Data governance* and internal *policy governance* are most strongly associated with AI governance by those that practice it, based on the same survey. A well-rounded AI governance approach must address technology architecture, data, models, processes and people.

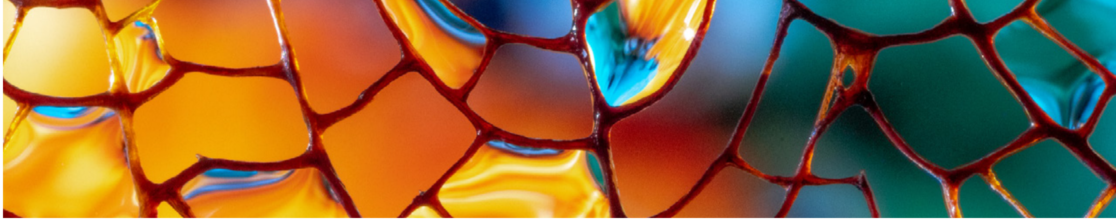

Business Impact

AI and GenAI can positively transform business, but also introduce significant risk. GenAI promises to augment human creativity and productivity, yet it requires access to potentially sensitive business data sources. Organizations must adequately protect this data to avoid incurring cyber and competitive risks.

Protecting sensitive data and IP is critical to the safe and responsible use of AI and GenAI. In the organizational use of AI and GenAI, it is rarely the model that determines differentiating outcomes. LLMs are largely commoditized. Proprietary business datasets typically differentiate outcomes and hence are the “secret sauce” that needs to be protected. Customer data, personal data and IP all must be considered.

Privacy and security are important, but organizations must also consider choice, cost, compliance and performance. Data security and privacy may be “solved” by locking everything down to the detriment of end users and use cases, but rigid methodologies can hamstring the business. An architectural approach needs to balance privacy, security and compliance concerns with considerations such as performance.

Poor privacy can be costly for organizational AI and GenAI efforts, while robust privacy can accelerate outcomes. Insufficient privacy practices can expose sensitive data to adversarial attacks, inadvertent model training and inappropriate access or use. Robust privacy — when the business balances control with usability and access — can foster a safe environment for AI experimentation and growth.

AI governance can bolster privacy and compliance outcomes, but compliance should not be the only objective. Organizations have historically taken a reactive “checkbox” approach to compliance, but trying to continually catch up with evolving regulations is a losing game. AI governance, done correctly, is proactive and supports the scalability and repeatability of AI efforts.

Looking ahead

Organizations are fixated on the business potential for AI and GenAI. Customer support, personalization, code generation, fraud detection, scientific simulation and knowledge summarization are all viable use cases, and new ones are sure to be uncovered as the technology matures. Yet as global regulations for AI continue to evolve — such as the EU AI Act — businesses will need to exercise caution as they adopt and use AI. There is enormous competitive pressure to keep up and innovate. To ignore AI for the sake of compliance or security is not viable, but at the same time, external forces are demanding more transparency of AI methodology as well as more robust protection of sensitive and personal data. Savvy organizations will raise their internal standards for the governance, security, privacy and use of AI before regulators come knocking.

In today’s complex AI technology ecosystem, it can be difficult for organizations to make sense of privacy, security and safety choices for their platforms and data. Yet while we often speak of AI and GenAI in the context of risk, the technology can also accelerate risk mitigation. GenAI is particularly deft at making sense of unstructured data, which historically has been poorly governed and categorized. The leverage of GenAI, using privacy-protecting methods, could accelerate organizations’ control and security of this information. Control and security of the enterprise AI ecosystem need to be a priority. To reap the rewards of emergent technology, organizations need to have robust guardrails in place.

vmware®
by **Broadcom**

Broadcom and NVIDIA aim to solve these issues and unlock the power of GenAI and unleash productivity with a joint GenAI platform – VMware Private AI Foundation with NVIDIA. This platform is built and run on the private cloud platform, VMware Cloud Foundation, and comprises the new NVIDIA NIM inference microservices, AI models from NVIDIA and others in the community (such as Hugging Face), and NVIDIA AI tools and frameworks, which are available with NVIDIA AI Enterprise licenses. This joint GenAI platform enables enterprises to run RAG workflows, fine-tune and customize LLM models, and run inference workloads in their data centers, addressing privacy, choice, cost, performance, and compliance concerns.

For more information about this platform, please visit this webpage.